14.12 Game Theory
Muhamet Yildiz
Fall 2005

Solution to Homework 1

Solution to Problem 1

a) (ten points) Let's first find appropriate payoffs for Alice. The information stated in the question shows us that the payoffs will be of the form

		Bob		
		A	M	H
Alice	P	a	b	b
	G	b	a	c

We can arbitrarily assign b to be 0 and (using the information that b>c) c to be -1. All that remains is to solve for a. To do this, we'll use the information that Alice prefers Penn Station to Grand Central Station iff p>q-r/2. In other words, we have

$$ap + 0q + 0(1 - p - q) > 0p + aq + (-1)(1 - p - q)$$

iff $p > q - (1 - p - q)/2$. We can rearrange these inequalities to see that we want $a(p - q) > (-1)(1 - p - q)$ iff $p - q > -(1 - p - q)/2$. In other words, we want $a(p - q) > (-1)(1 - p - q)$ iff $2(p - q) > -(1)(1 - p - q)$. So we see that we need $a = 2$.

Now let's look at Bob's payoffs. The information stated in the question shows us that they must be of the form

		Bob		
		A	M	H
Alice	P	x	z	w
	G	z	y	w

Moreover, when the probability of Alice waiting at Penn Station is 1/2, Bob prefers Amtrak to Metroliner, which tells us that x>y. Using this information, we can arbitrarily assign x to be 1 and y to be 0. Now we need to solve for z and w. Using s to denote the probability of Alice waiting at Penn Station, the fact that Bob prefers Amtrak to Metroliner iff s $> 1/3$ means that $1s + z(1 - s) > zs + 0(1 - s)$ iff s>1/3. This will hold if z=-1. Finally, the fact that Bob prefers Amtrak to Home iff s>2/3 means that $1s + (-1)(1 - s) > w$ iff s>2/3. This means that w=1/3.

Thus, in the end, our normal form game is:

		Bob		
		A	M	H
Alice	P	2,1	0,-1	0,1/3
	G	0,-1	2,0	-1,1/3

b) (five points) We can perform any positive affine transformation to the payoff function of either player and preserve the structure of the game. So let's add 1 to Alice's payoffs and multiply Bob's payoffs by 3. We end up with:

		Bob		
		A	M	H
Alice	P	3,3	1,-3	1,1
	G	1,-3	3,0	0,1

c) (ten points) The question is asking for the rationalizable strategy profiles. Note that H strictly dominates M. Removing M from the game gives us:

		Bob	
		A	H
Alice	P	3,3	1,1
	G	1,-3	0,1

In this new game above, P strictly dominates G. Removing G gives us:

		Bob	
		A	H
Alice	P	3,3	1,1

And in the game where Alice goes to Penn Station and Bob chooses between A and H, A strictly dominates H. So the only possible outcome given that it's common knowledge that both players are expected utility maximizers with the stated preferences is for Bob to take Amtrak and for Alice to meet him at Penn Station.

Solution to Problem 2

a) (six points)

		2							
		Llλ	Llρ	Lrλ	Lrρ	Rlλ	Rlρ	Rrλ	Rrρ
	Axa	3,0	3,0	3,0	3,0	0,2	0,2	0,2	0,2
	Axb	3,0	3,0	3,0	3,0	0,2	0,2	0,2	0,2
	Aya	1,5	1,5	1,5	1,5	0,2	0,2	0,2	0,2
	Ayb	1,5	1,5	1,5	1,5	0,2	0,2	0,2	0,2
	Bxa	1,2	1,2	2,1	2,1	1,2	1,2	2,1	2,1
1	Bxb	1,2	1,2	2,1	2,1	1,2	1,2	2,1	2,1
	Bya	1,2	1,2	2,1	2,1	1,2	1,2	2,1	2,1
	Byb	1,2	1,2	2,1	2,1	1,2	1,2	2,1	2,1
	Cxa	1,0	2,1	1,0	2,1	1,0	2,1	1,0	2,1
	Cxb	0,1	0,0	0,1	0,0	0,1	0,0	0,1	0,0
	Cya	1,0	2,1	1,0	2,1	1,0	2,1	1,0	2,1
	Cyb	0,1	0,0	0,1	0,0	0,1	0,0	0,1	0,0

b) (seven points) For player 1, Cxa (among other things) strictly dominates Cxb and Cyb. Moreover, the mixed strategy that involves playing Axa and Bxa each with probability 1/2 strictly dominates Aya, and it also strictly dominates Ayb. Every other strategy of player 1 is a best response to some strategy of player 2, which implies that it can't be strictly dominated. Similarly, each strategy of player 2 is a best response to some strategy of player 1, so player 2 has no strictly dominated strategies. After deleting strictly dominated strategies, we end up with:

					2				
		L lλ	Llρ	Lrλ	Lrρ	Rlλ	Rlρ	Rrλ	Rrρ
	Axa	3,0	3,0	3,0	3,0	0,2	0,2	0,2	0,2
	Axb	3,0	3,0	3,0	3,0	0,2	0,2	0,2	0,2
	Bxa	1,2	1,2	2,1	2,1	1,2	1,2	2,1	2,1
1	Bxb	1,2	1,2	2,1	2,1	1,2	1,2	2,1	2,1
	Bya	1,2	1,2	2,1	2,1	1,2	1,2	2,1	2,1
	Byb	1,2	1,2	2,1	2,1	1,2	1,2	2,1	2,1
	Cxa	1,0	2,1	1,0	2,1	1,0	2,1	1,0	2,1
	Cya	1,0	2,1	1,0	2,1	1,0	2,1	1,0	2,1

In this new 8 x 8 game, Lrλ is strictly dominated by a mixed strategy that plays Llρ and Rrλ each with probability 1/2, but every other strategy for each player is a best response to some strategy of the other player. So after the second round of elimination, we have:

					2			
		Llλ	Llρ	Lrρ	Rlλ	Rlρ	Rrλ	Rrρ
	Axa	3,0	3,0	3,0	0,2	0,2	0,2	0,2
	Axb	3,0	3,0	3,0	0,2	0,2	0,2	0,2
	Bxa	1,2	1,2	2,1	1,2	1,2	2,1	2,1
1	Bxb	1,2	1,2	2,1	1,2	1,2	2,1	2,1
	Bya	1,2	1,2	2,1	1,2	1,2	2,1	2,1
	Byb	1,2	1,2	2,1	1,2	1,2	2,1	2,1
	Cxa	1,0	2,1	2,1	1,0	2,1	1,0	2,1
	Cya	1,0	2,1	2,1	1,0	2,1	1,0	2,1

In this new 8 x 7 game, each strategy of either player here is a best response to some strategy of the opponent in this game, which implies that we can no longer eliminate anything. So the rationalizable strategies are Axa, Axb, Bxa, Bxb, Bya, Byb, Cxa, and Cya for player 1; and Llλ, Llρ, Lrρ, Rlλ, Rlρ, Rrλ, and Rrρ for player 2.

c) (five points) The best responses of player 1 to Llλ are Axa and Axb. Player 1's payoff from playing those strategies are in bold in the column corresponding to Llλ in the payoff matrix below. I have done this for player 1's best response to each of the rest of player 2's strategies, and I have done a similar thing for player 2's best response to each of player 1's strategies. If both payoffs in a box are in bold, it means that each player is playing a best response to the other (i.e., we have a Nash equilibrium).

2

		Llλ	Llρ	Lrλ	Lrρ	Rlλ	Rlρ	Rrλ	Rrρ
	Axa	**3**,0	**3**,0	**3**,0	**3**,0	0,**2**	0,**2**	0,**2**	0,**2**
	Axb	**3**,0	**3**,0	**3**,0	**3**,0	0,**2**	0,**2**	0,**2**	0,**2**
	Aya	1,**5**	1,**5**	1,**5**	1,**5**	0,2	0,2	0,2	0,2
	Ayb	1,**5**	1,**5**	1,**5**	1,**5**	0,2	0,2	0,2	0,2
	Bxa	1,**2**	1,**2**	2,1	2,1	**1**,**2**	1,**2**	**2**,1	**2**,1
1	Bxb	1,**2**	1,**2**	2,1	2,1	**1**,**2**	1,**2**	**2**,1	**2**,1
	Bya	1,**2**	1,**2**	2,1	2,1	**1**,**2**	1,**2**	**2**,1	**2**,1
	Byb	1,**2**	1,**2**	2,1	2,1	**1**,**2**	1,**2**	**2**,1	**2**,1
	Cxa	1,0	2,**1**	1,0	2,**1**	**1**,0	**2**,**1**	1,0	**2**,**1**
	Cxb	0,**1**	0,0	0,**1**	0,0	0,**1**	0,0	0,**1**	0,0
	Cya	1,0	2,**1**	1,0	2,**1**	**1**,0	**2**,**1**	1,0	**2**,**1**
	Cyb	0,**1**	0,0	0,**1**	0,0	0,**1**	0,0	0,**1**	0,0

So there are eight Nash equilibria in pure strategies: (Bxa,Rlλ), (Bxb,Rlλ), (Bya,Rlλ), (Byb,Rlλ), (Cxa,Rlρ), (Cxa,Rrρ), (Cya,Rlρ), and (Cya,Rrρ).

d) (seven points) For player 1, Aya, Ayb, Cxb, and Cyb are weakly dominated by Axa. None of the other strategies are dominated (not even by a mixed strategy) because they are each a best response to some strategy of player 2 and there is no other strategy of player 1 that is also a best response and that gives a higher payoff against any of player 2's strategies. For player 2, Rlρ weakly dominates Rrρ, Rlλ weakly dominates Rrλ, Llλ weakly dominates Lrλ, and Llρ weakly dominates Lrρ. But neither Rlρ,Rrλ, Llλ, nor Llρ is weakly dominated (you should be able to convince yourself of this). So after round one of elimination, we end up with:

2

		Llλ	Llρ	Rlλ	Rlρ
	Axa	3,0	3,0	0,2	0,2
	Axb	3,0	3,0	0,2	0,2
	Bxa	1,2	1,2	1,2	1,2
1	Bxb	1,2	1,2	1,2	1,2
	Bya	1,2	1,2	1,2	1,2
	Byb	1,2	1,2	1,2	1,2
	Cxa	1,0	2,1	1,0	2,1
	Cya	1,0	2,1	1,0	2,1

Now for player 1, Cxa weakly dominates Bxa, Bxb, Bya, and Byb. Axa, Axb, Cxa, and Cya are not weakly dominated. For player 2, Rlρ is a weakly dominant strategy. Thus, after round two of elimination, we have:

		2
		Rlρ
	Axa	0,2
1	Axb	0,2
	Cxa	2,1
	Cya	2,1

And in this game, Cxa weakly dominates Axa and Axb. So after round three, we wind up with:

		2
		Rlρ
1	Cxa	2,1 .
	Cya	2,1

And it should be obvious that we can do no more elimination. So the strategies that survive iterated elimination of weakly dominated strategies are Cxa and Cya for player 1, and Rlρ for player 2.

Solution to Problem 3.

(a) (10 points) Payoff to student i is

$$x_i - x_i t\left(x_1, .., x_n\right)$$

Given the choices of the other students $\left\{x_j^*\right\}_{j \neq i}$, player i's optimal strategy is to choose x_i such that

$$1 - x_i \frac{\partial t}{\partial x_i} - t\left(x_1^*, x_2^* .. x_i, .., x_n^*\right) = 0$$

$$\Longrightarrow x_i = 1 - t\left(x_1^*, x_2^* .. x_i, .., x_n^*\right) \tag{1}$$

Summing across all students,

$$\sum_{j=1}^{n} x_j^* = t\left(x_1^*, .., x_n^*\right) = n - nt\left(x_1^*, .., x_n^*\right)$$

$$\Longrightarrow t\left(x_1^*, .., x_n^*\right) = \frac{n}{n+1} \tag{2}$$

Substituting for $t\left(x_1^*, .., x_n^*\right)$ in (1) using (2),

$$x_i = \frac{1}{n+1}$$

Therefore, in the unique Nash equilbrium, each student chooses to send data of size $\frac{1}{n+1}$ and receives a payoff of $\frac{1}{(n+1)^2}$.

(b) (15 points)

1. Payoff to student i is

$$M + x_i - x_i t(x_1, .., x_n) - px_i$$

We proceed as in the first part. If the choices of the other students are given by $\{x_j^*\}_{j \neq i}$, player i's optimal strategy is given by

$$1 - x_i \frac{\partial t}{\partial x_i} - t(x_1^*, x_2^* .. x_i, .., x_n^*) - p = 0$$
$$\implies x_i = 1 - t(x_1^*, x_2^* .. x_i, .., x_n^*) - p \qquad (3)$$

Summing across all students,

$$\sum_{j=1}^{n} x_j^* = t(x_1^*, .., x_n^*) = n - nt(x_1^*, .., x_n^*) - np$$

$$\implies \quad t(x_1^*, .., x_n^*) = \frac{n(1-p)}{n+1}$$
$$\implies \quad x_i^* = 1 - \frac{n(1-p)}{n+1} - p = \frac{1-p}{n+1}$$

Therefore, given M and p, each player chooses to send $\frac{1-p}{n+1}$ units of data in the Nash equilibrium, and receives a payoff equal to $M + \left[\frac{1-p}{n+1}\right]^2$

2. If each student sends $\frac{1-p}{n+1}$ units of data, at price p this generates a total revenue of $\frac{np(1-p)}{n+1}$. So, to break even, we must set

$$nM = \frac{np(1-p)}{n+1} \implies M = \frac{p(1-p)}{n+1}$$

The utility each player obtains from such a scheme =

$$\frac{p(1-p)}{n+1} + \left[\frac{1-p}{n+1}\right]^2$$
$$= \frac{1}{(n+1)^2}(1-p)(np+1)$$

which attains its maximum at $p = \frac{n-1}{2n}$, $M = \frac{n-1}{4n^2}$. These values for p and M yield a utility of $\frac{1}{4n}$.

3. Note that, for $n = 1$, the expressions obtained for p and M in part (b)2 equal zero and the utility to each player in equilibrium equals $\frac{1}{4}$ in both the first and second programs. The reason this is so is that when there is only one individual, there is no scope for negative externalities on the data network, and therefore the charge for

transmitting data cannot improve utility. For higher values of n, the second program always does better – by charging students for data transmission, it ensures that they take into account the negative externalities involved when deciding how much data to send.

Solution to Problem 4

The following is a proof for the general n-person case [a proof of the 3-person case previously posted contained an error and has been removed; it will be substituted with an alternate proof]. Denote by t_i the true preferences of candidate i. Fix a set of 'declared' prefereces $\{s_j\}_{j\neq i}$ for the other candidates (which may or may not involve truth-telling). Denote by μ_M the matching that takes place when the algorithm described in the question (called henceforth the Gale-Shapley algorithm) is used for declared preferences $\left(t_i, \{s_j\}_{j\neq i}\right)$. We show below that if there is any other matching μ such that candidate i prefers his position under μ to that under μ_M, then the matching μ will be 'blocked'; i.e. there is a candidate (different from i) and a position not paired with one another under μ who prefer such a pairing (according to their candidate's declared preferences) to their own partners under μ. In the Gale-Shapley paper, they show that their matching algorithm will always lead to an outcome that is 'stable' (for the declared preferences); i.e. there is no possible blocking pair. Therefore, if candidate i prefers μ over μ_M, it is not possible to obtain μ using the Gale-Shapley algorithm for any preference s_i declared by candidate i.

To show that that matching μ will be 'blocked' under preferences $\left(t_i, \{s_j\}_{j\neq i}\right)$: Let M' be the group of candidates who prefer μ to μ_M; by assumption M' includes candidate i but it may also include other candidates. Denote by $\mu(M')$ and $\mu_M(M')$ the group of positions to which the candidates in M' are matched under μ and μ_M respectively. There are two cases that we need to consider separately (the following is based on the the proof of the 'Blocking Lemma' from 'Two-sided Matching' by Roth and Sotomayer).

- Case 1: $\mu(M') \neq \mu_M(M')$. Choose w in $\mu(M') - \mu_M(M')$. That is, w is a position matched to a candidate in M' under μ but not under μ_M. Say $w = \mu(m')$. Because $m' \in M'$, m' prefers w to the position to which he was matched under μ_M. Then m' must have applied to w under μ_M and been rejected; so it must be that w prefers $\mu_M(w) = m$ to m'. Since $w \notin \mu_M(M')$, $m \notin M'$. So m prefers w to $\mu(m)$. So (m, w) 'blocks' μ.

- Case 2: $\mu_M(M') = \mu(M') = W'$. Let w be the last position in W' to receive an application from a member of M' under μ_M. Since all the positions in W' have rejected candidates in M', it must be that w had some applicant m engaged when it received this last application.Then, we claim (m, w) is the blocking pair. First, $m \notin M'$; otherwise he would apply to some other position in W' after being rejected by w, and this would contradict the assumption that w is the last position in W' to receive such an application. So m is worse off under μ than under μ_M. Also, since

m applies to w (and is rejected) before being matched with $\mu_M(m)$, he prefers w to $\mu_M(m)$. Therefore, he also prefers w to $\mu(m)$. On the other hand, m was the last candidate to be rejected by w; therefore, she must have rejected $\mu(w)$ before she rejected m. Therefore, she prefers m to $\mu(w)$. So (m, w) blocks μ as claimed.